arXiv:1701.08280v2 [cs.CV] 16 Feb 2017

Pruned Non-Local Means

Sanjay Ghosh, Amit K. Mandal, and Kunal N. Chaudhury

Abstract

In Non-Local Means (NLM), each pixel is denoised by performing a weighted averaging of its neighboring pixels, where the weights are computed using image patches. We demonstrate that the denoising performance of NLM can be improved by pruning the neighboring pixels, namely, by rejecting neighboring pixels whose weights are below a certain threshold λ. While pruning can potentially reduce pixel averaging in uniform-intensity regions, we demonstrate that there is generally an overall improvement in the denoising performance. In particular, the improvement comes from pixels situated close to edges and corners. The success of the proposed method strongly depends on the choice of the global threshold λ, which in turn depends on the noise level and the image characteristics. We show how Stein's unbiased estimator of the mean-squared error can be used to optimally tune λ, at a marginal computational overhead. We present some representative denoising results to demonstrate the superior performance of the proposed method over NLM and its variants.

I. Introduction

We consider the standard problem of denoising grayscale images that are corrupted with additive white Gaussian noise. The setting here is that we are given a *noisy* image

$$y_i = x_i + n_i \qquad (i \in I), \tag{1}$$

where I is some finite rectangular domain of $\mathbb{Z}^2$, $\{x_i : i \in I\}$ is the unknown *clean* image, and $\{n_i : i \in I\}$ are i.i.d. $\mathcal{N}(0, \sigma^2)$ where σ is the standard deviation of the Gaussian noise. The focus of this work is the Non-Local Means (NLM) algorithm that was introduced by Buades, Coll, and Morel in [1], [2]. NLM denoises each pixel by taking a weighted average of the neighboring pixels. A key innovation in NLM is that the weights are computed using patches (blocks of neighboring pixels) instead of pixels. More specifically, the *denoised* image $\{\hat{x}_i : i \in I\}$ is computed using the formula

$$\hat{x}_i = \frac{\sum_{j \in S(i)} w_{i,j} y_j}{\sum_{j \in S(i)} w_{i,j}}, \tag{2}$$

Correspondence: kunal@ee.iisc.ernet.in. Department of Electrical Engineering, Indian Institute of Science.

where $S(i)$ denotes the *neighborhood* around pixel i. The original proposal in [1] was to set $S(i)$ to be the whole image. In practice, one restricts $S(i)$ to a sufficiently large window of size $(2S+1)\times(2S+1)$ centered at i, that is, $S(i) = i + [-S, S]^2$. The weights $\{w_{i,j} : i \in I,\ j \in S(i)\}$ are set to be

$$w_{i,j} = \exp\Big(-\frac{1}{h^2}\sum_{k\in\mathcal{P}}(y_{i+k} - y_{j+k})^2\Big), \tag{3}$$

where h is a smoothing parameter, and *patch* $\mathcal{P} = [-K, K]^2$ is a square neighborhood of the origin of size $(2K + 1) \times (2K + 1)$.

Since its introduction, several variations of NLM have been proposed which can further improve the denoising. For example, various weighting schemes for the center pixel have been proposed in [3], [4], [5]. A weighting scheme based on quadratic programming was proposed in [6], which exploits the overlapping information between adjacent patches. Further, optimization framework based approaches have been proposed in [7], [8] to improve the performance of NLM. In [9], the weighting was performed based on the structural similarity indices between patches. In a different direction, it was demonstrated in [10], [11], [12], [13] that the performance (and the speed) of NLM can be improved by computing the weights from low-dimensional projections of the patches.

The primary application of NLM is noise reduction [14], [15]. However, since the inception of NLM, more sophisticated algorithms with superior denoising performances have been proposed [16], [17], [18], [19], [20], [21], [22], [23], [24]. While NLM and its variants are no longer the state-of-the-art in image denoising, they nevertheless continue to be of interest due to their simplicity, decent denoising performance, and the availability of fast implementations. We also note that several researchers have used NLM (either directly or via post-processing) in image restoration applications such as deblurring [25], [26], inpainting [27], demosaicking [28], and super-resolution [29], [30], [31].

The investigation in this paper was particularly motivated by the work on shape-adaptive patches [32] and robust ℓ_1 regression [33], [34]. In either case, it was demonstrated that the denoising quality of NLM can be improved by adapting the averaging neighborhood at every pixel. While this was done explicitly in [32] in terms of the patch shape, the adaptation was achieved implicitly in [33], [34] via an optimization. In this paper, we propose to perform the adaptation by explicitly pruning those neighboring pixels whose weights are below a certain global threshold λ. We use Stein's unbiased estimator of the mean-squared error to optimally tune λ. In other words, we optimize the denoising quality by performing a simple one-dimensional optimization for the whole image. As a result, the overall implementation is much faster than those implementations in [32], [33], [34], [35], where one is required to perform an adaptation at every pixel.

The rest of the paper is organized as follows. In Section II, we motivate the main idea using a simple experiment, and then present the proposed denoising algorithm and the implementation details.

The denoising results are reported in Section III, where we also compare the proposed method with NLM and its variants.

II. Pruned Non-Local Means

We now perform some numerical experiments to reflect the motivation that pruning can improve the denoising performance. We consider a scenario where the pixel of interest (POI) is situated closed to a edge, as shown in the example in Figure 1(a). We compute the NLM weights around the POI (marked with a solid black circle) using (3). We expect the weights of neighboring pixels that are on the same side of the edge as the POI to be generally higher than the weights of pixels that are on the other side of the edge. As shown in Figure 1(b) this is indeed the case if the weights are computed from the (unknown) clean edge. Unfortunately, this is not the case if the weights are computed from the noisy edge. As seen in Figure 1(b), the weights of some of the pixels on the left of the edge (same side of the edge as the POI) are smaller than the weights of some of the pixels on the right. This leads to a poor averaging in (2) – pixels from both sides of the edge are mixed. For example, in Figure 1(a), the intensities at the POI before and after adding noise are 0 and -3.75. Using NLM, the denoised output is 81.72 (as a reference, the NLM output using the weights computed from the clean patches is 73.09). On the other hand, notice that if we sort the noisy weights in Figure 1(b), then the majority of the pixels on the left of the edge would have the largest weights. Indeed, if we use the top 50% neighbors (neighbors with largest weights) for the averaging, then the output comes down to 42.32 (which is even better than that obtained using the clean weights). This suggests a means of improving the quality of denoising, namely, by restricting the averaging in (2) to those neighboring pixels that are assigned top weights. Based on the above observation, we propose to perform *pruning* as follows. We fix a threshold $0 < \lambda < 1$, and consider the step function $\phi_\lambda : [0,1] \to [0,1]$ given by

$$\phi_\lambda(t) = \begin{cases} 1, & \text{for } \lambda \leq t \leq 1 \\ 0, & \text{for } 0 \leq t < \lambda. \end{cases} \tag{4}$$

The process of nullifying the weights that are smaller than λ can be achieved by modifying (1) as:

$$\hat{x}_i = \frac{\sum_{j \in S(i)} \; w_{i,j} \; \phi_\lambda(w_{i,j}) \; y_j}{\sum_{j \in S(i)} w_{i,j} \; \phi_\lambda(w_{i,j})}. \tag{5}$$

To avoid introducing unnecessary symbols, we continue to use the same notation for the denoised image as in (2). Notice that the weights defined in (3) take values in $[0,1]$; therefore, the variable $t \in [0,1]$ in (4). When $\lambda = 0$, none of the weights are pruned, and when λ is close to unity, a large fraction of the weights are pruned. The optimal setting would be somewhere in between. We can ideally tune λ by minimizing the mean-squared-error (MSE) given by

$$\text{MSE} = \frac{1}{|I|} \sum_{i \in I} \big(\hat{x}_i - x_i\big)^2, \tag{6}$$

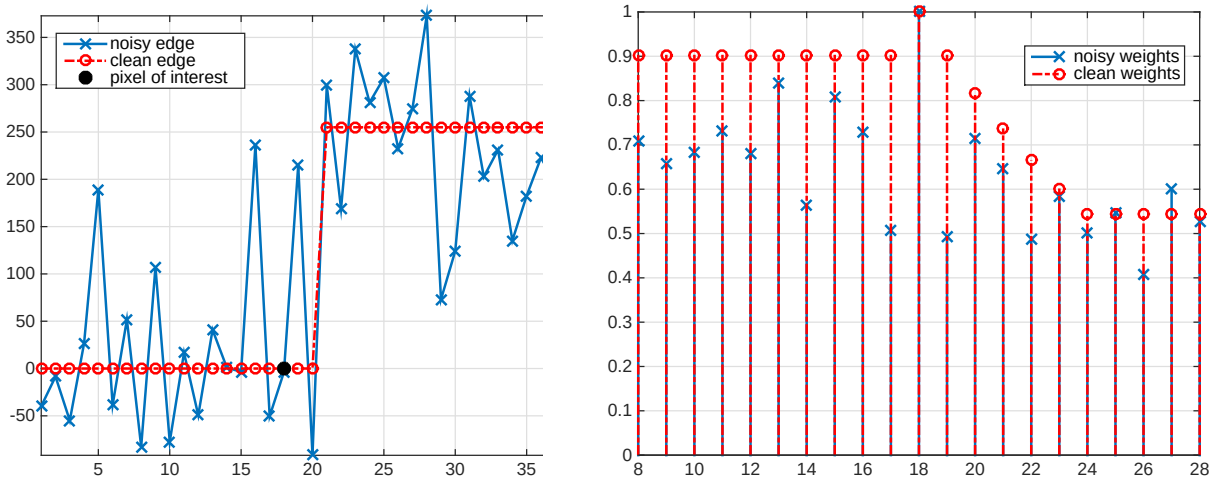

(a) clean and noisy edges ($\sigma = 80$). (b) neighboring weights for the POI.

Fig. 1. NLM-based denoising at a pixel (index = 18) situated close to an ideal edge. NLM parameters: $S = 10, K = 3$, and $h = 10\sigma$.

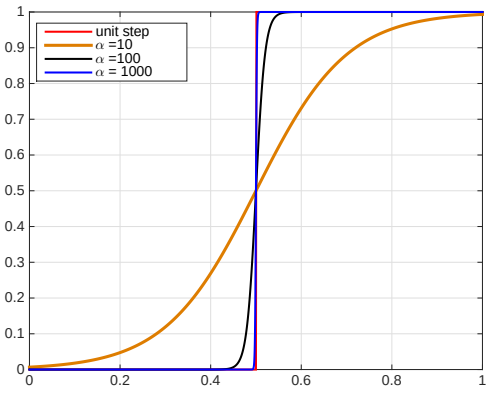

Fig. 2. Sigmoid function in (8) when $\lambda = 0.5$. At large values of α, we obtain a close approximation of the unit step function.

where $|I|$ is the size of the image, i.e., the number of pixels in the image. However, since we do not have access to the clean image $\{x_i\}$, we cannot compute the MSE in practice. This is precisely where Stein's Unbiased Risk Estimator (SURE) proves to be useful. This is given by

$$\text{SURE} = \frac{1}{|I|}\sum_{i\in I}\left(\hat{x}_i - y_i\right)^2 - \sigma^2 + \frac{2\sigma^2}{|I|}\sum_{i\in I}\frac{\partial \hat{x}_i}{\partial y_i}, \tag{7}$$

where we recall that σ is the standard deviation of the additive Gaussian noise in (1). SURE has the remarkable property that its expected value equals that of the MSE for the noise model in (1) [36]. This makes it an useful surrogate for the MSE – it can be computed without the knowledge of the clean image. The difficulty in applying SURE directly to the transformation in (5) is that (4) is not differentiable and therefore we cannot compute the partial derivatives in (7). This is, however, not

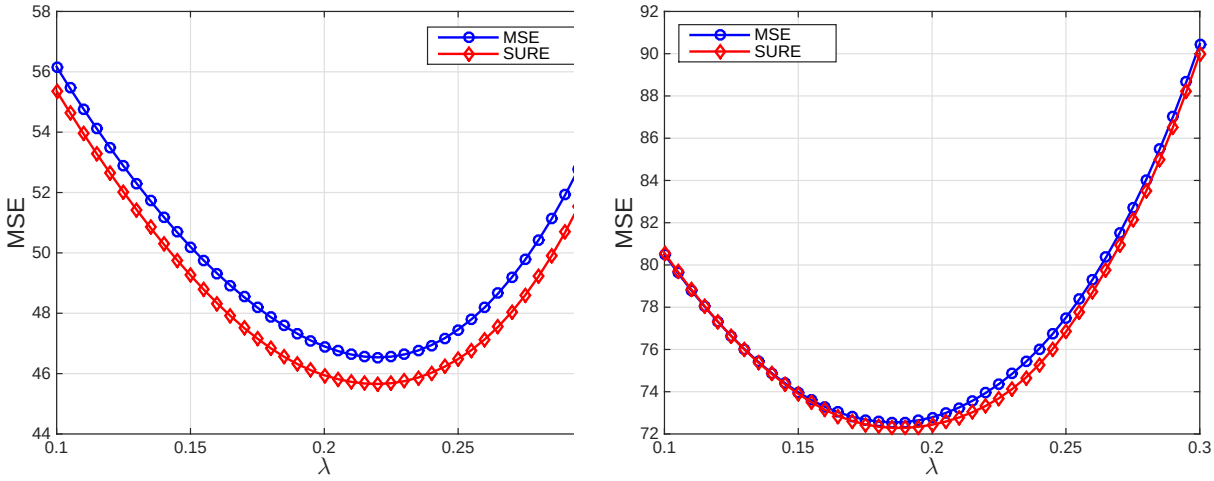

(a) clean and noisy edges ($\sigma = 80$). (b) clean and noisy edges ($\sigma = 80$).

Fig. 3. Comparison of MSE (computed from the clean image) and SURE as a function of λ for two different images ($\sigma = 20$).

a fundamental problem and we can approximate (4) by a differentiable function. For example, we propose to use the following sigmoid function:

$$\phi_{\lambda,\alpha}(t) = \frac{1}{1 + e^{-\alpha(t-\lambda)}} \qquad (0 \leq t \leq 1), \tag{8}$$

where $\alpha > 0$ is the slope of the sigmoid at $t = \lambda$. It can be verified that (8) converges pointwise to (4) as α get large, namely, for fixed $0 \leq t \leq 1$,

$$\lim_{\alpha \to \infty} \phi_{\lambda,\alpha}(t) = \phi_{\lambda}(t).$$

In Figure 2, we plot (8) for different values of α and compare it with (4) for the same λ. Ideally, we should use a large α to obtain a good approximation. However, as will be evident shortly, the partial derivatives in (7) explode as α increases. Following exhaustive simulations, we decided to set $\alpha = 100$. The approximation is reasonable in this case and the derivatives can be computed stably. In other words, we consider the following formula

$$\hat{x}_i = \frac{\sum_{j \in S(i)} w_{i,j} \phi_{\lambda,\alpha}(w_{i,j})\ y_j}{\sum_{j \in S(i)} w_{i,j} \phi_{\lambda,\alpha}(w_{i,j})}, \tag{9}$$

where $\alpha = 100$. We will refer to (9) as Pruned Non-Local Means (PNLM). We note that the SURE for (2) was originally derived in [37], and was shown to be close to the MSE. Similarly, we can derive the SURE for (9). In particular, the partial derivative of (9) is given by (cf. Appendix):

$$\frac{\partial \hat{x}_i}{\partial y_i} = \frac{1}{W_i}\Big[\psi(1) + \frac{2}{h^2}\sum_{j \in S(i)} w_{i,j}\ \psi'(w_{i,j})(y_j - \hat{x}_i)(y_j - y_i)$$
$$+ \frac{2}{h^2}\sum_{k \in \mathcal{P}} w_{i,i+k}\ \psi'(w_{i,i+k})(y_{i+k} - \hat{x}_i)(y_{i-k} - y_i)\Big], \tag{10}$$

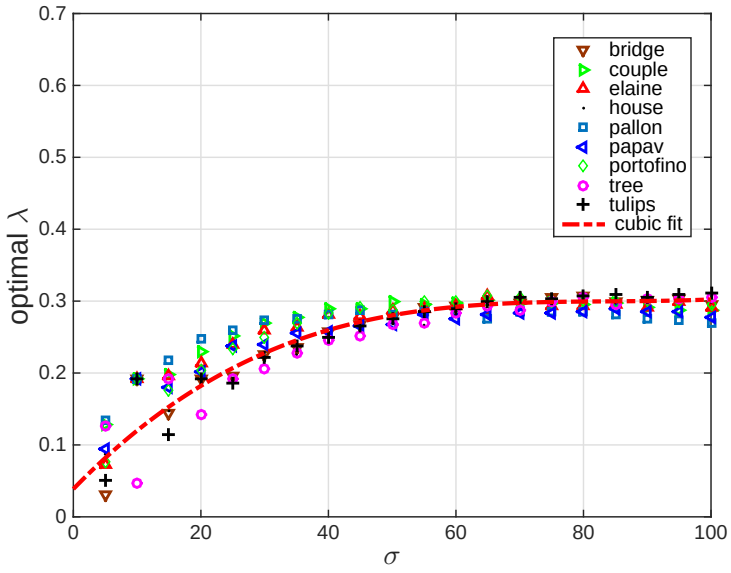

Fig. 4. Optimal λ versus noise level for various natural images [41], [42]. The parameters used are $S = 10$, $K = 3$, and $h = 10\sigma$.

where

$$W_i = \sum_{j \in S(i)} \psi(w_{i,j}) \quad \text{and} \quad \psi(t) = t\phi_{\lambda,\alpha}(t). \tag{11}$$

The expression of $\psi'(t)$ is given in (17) for completeness. A comparison of the MSE and SURE is provided in Figure 3 when $\alpha = 100$. We notice that the gap between MSE and SURE decreases with the increase in image size. Importantly, we note that the minima of MSE and SURE are attained at almost the same λ.

III. Experimental Results

A. Parameters

The denoising performance of PNLM clearly depends on the choice of λ. In particular, the optimal λ^* (minimizer of MSE) depends on the noise level and the image characteristics. In order to obtain a noise-based rule for setting λ^*, we took a large collection of natural images [41], [42] and computed the optimal λ at various noise levels using exhaustive search. The scatter plot of σ and the corresponding λ^* for some of the images is shown in Figure 4. Also shown is the cubic polynomial that we obtained using least-square fitting:

$$\lambda^* = 4.3 \times 10^{-7}\sigma^3 - 1.1 \times 10^{-4}\sigma^2 + 9.2 \times 10^{-3}\sigma + 0.039. \tag{12}$$

We note that the above fitting pertains to the parameter setting $S = 10, K = 3$ and $h = 10\sigma$. In this case, we can directly use (12) to predict λ^*. Of course, (12) will not provide a good estimate of λ^* for other parameter settings. One option is to maintain a lookup table of the optimal rules for selected

parameter settings, and then interpolate between them. Instead, we propose to use (12) to initialize a one-dimensional search, irrespective of the parameter setting. Starting from a coarse estimate of λ^*, we then use golden section search [38] to refine the estimate. The complete process is summarized in Algorithm 1. Note that, compared to NLM, we have to perform an additional search in PNLM. Each step of this search requires us to compute (7) and (9) twice. However, the weights in (7) and (9) are computed using the patch distances used in NLM. Therefore, we have to compute the patch distances just once, which happens to be the computation-intensive part of NLM. The iterations in Algorithm 1 are terminated when the change in λ^* between successive iterations is within 10^{-4}. The whole process is quite fast once the patch distances have been computed.

The denoising performance of PNLM and NLM are compared for different choices of search window and patch size in Table I. We see that PNLM achieves better denoising irrespective of the choice of parameters. We note that the optimal selection of NLM parameters is an altogether different line of work [3], [12]. In this paper, the focus is on post-processing of weights to enhance the robustness of NLM. The idea of pruning remain relevant irrespective of the choice of NLM parameters. For the present discussion, we choose the NLM parameters to be $S = 10, K = 3$ and $h = 10\sigma$, as was originally suggested in [1]. As far as the additional parameters α and λ in PNLM are concerned, we set $\alpha = 100$ to ensure a reasonable approximation of the step function ϕ_λ, and the crucial parameter λ is set using the process described above.

Algorithm 1 Golden Section Search

Input: Noisy image $\{y_i\}$; noise level σ; parameters S, K, and h.

Return: Estimate of λ^*.

Initialize: $k = 0$; $\rho = 0.618$; $\delta = 0.05$; set λ_0 using (12).

Loop until converge

1. Set $l_k = \lambda_k - \delta$ and $u_k = \lambda_k + \delta$.

2. Set $p_k = u_k - \rho(u_k - l_k)$ and $q_k = l_k + \rho(u_k - l_k)$.

3. If $\mathrm{SURE}(p_k) > \mathrm{SURE}(q_k)$

   set $l_{k+1} = p_k$ and $u_{k+1} = u_k$.

   Else

   set $l_{k+1} = l_k$ and $u_{k+1} = q_k$.

   $k = k + 1$;

End Loop

return $(l_k + u_k)/2$.

TABLE I

COMPARISON OF THE DENOISING PERFORMANCE (PSNR/SSIM) FOR DIFFERENT SEARCH WINDOW (S) AND PATCH SIZE (K). NOISE LEVEL USED IS $\sigma = 20$. THE RESULTS ARE FOR THE **HOUSE** IMAGE [41].

K	S	NLM	PNLM
$K = 2$	7	28.99 / 80.47	32.22 / 84.47
	10	28.31 / 79.98	31.97 / 84.36
	12	28.03 / 79.86	31.79 / 84.25
	15	27.60 / 79.53	31.52 / 84.00
	20	26.97 / 78.97	31.16 / 83.67
$K = 3$	7	30.45 / 82.42	32.26 / 84.77
	10	29.78 / 81.84	32.20 / 84.97
	12	29.45 / 81.59	32.19 / 85.06
	15	28.98 / 81.08	32.14 / 85.09
	20	28.33 / 80.35	32.05 / 85.03
$K = 4$	7	31.61 / 83.93	31.59 / 84.27
	10	31.21 / 83.60	31.73 / 84.53
	12	30.94 / 83.38	31.77 / 84.64
	15	30.52 / 82.92	31.78 / 84.69
	20	29.86 / 82.16	31.75 / 84.61

B. Runtime and Denoising Results

We now present simulation results concerning the run time of the proposed algorithm. In Table II, we compare the computation costs of NLM and PNLM on various image sizes. We see that the computational overhead for PNLM is within 30% of that of NLM. In this context, we note that one can use any of the existing fast algorithms for NLM to speedup the PNLM computation [39], [40]. By using an optimized C implementation of any of these fast algorithms, one can bring down the run-time to within a second for a 512×512 image.

TABLE II

COMPARISON OF THE RUN TIMES OF NLM AND PNLM FOR VARIOUS IMAGE SIZES, USING THE STANDARD PARAMETERS $S = 10$ AND $K = 3$. THE COMPUTATIONS WERE PERFORMED USING MATLAB 8.4 ON A 3.4 GHZ INTEL QUAD-CORE MACHINE WITH 32 GB MEMORY.

Method	64×64	128×128	256×256	512×512
NLM	7.7 s	30 s	122 s	490 s
PNLM	9.7 s	39 s	157 s	630 s

We now report some representative denoising results. We compared PNLM with NLM and the

following variants of NLM: NLEM [33], PCA-NLM [11], Zero-CPW [4], One-CPW [32], and Max-CPW [3]. For all experiments, we used the following parameter settings: $S = 10$, $K = 3$ and $h = 10\sigma$ [1]. In Table III and IV, we report the PSNR and SSIM [43] obtained using various denoising algorithms on the *Boat* and *Hill* images at different noise values. We use the standard definition

$$\text{PSNR} = 10\log_{10}(255^2/\text{MSE}), \tag{13}$$

where MSE is given by (6).

TABLE III

COMPARISON OF PSNR AND SSIM WITH NLM AND ITS VARIANTS FOR THE **BOAT** IMAGE [41]. THE HIGHEST PSNRS AND SSIMS ARE MARKED IN BOLD.

σ	NLM [1]	NLEM [33]	PCA-NLM [11]	Zero-CPW [4]	One-CPW [32]	Max-CPW [3]	PNLM
5	35.07/96.29	35.04/96.89	34.94/96.07	28.88/85.37	29.84/86.34	29.28/85.80	**35.78/97.90**
10	30.78/88.95	30.96/91.11	30.71/88.74	28.83/85.50	29.91/86.88	29.23/85.95	**32.22/94.37**
20	26.73/77.96	27.43/83.13	26.67/77.79	28.45/85.58	27.83/88.72	28.83/86.11	**28.95/87.12**
30	24.72/70.66	25.58/77.60	24.70/70.55	27.04/83.37	19.92/65.45	27.11/83.70	**27.29/81.74**
40	23.64/65.68	24.60/73.47	23.62/65.60	23.87/74.36	16.12/48.78	23.29/72.74	**26.33/77.63**
50	22.96/62.13	23.99/70.09	22.95/62.06	20.03/58.08	14.15/40.68	19.36/56.86	**25.18/73.92**
60	22.48/59.45	23.48/66.97	22.47/59.40	16.91/43.81	12.57/34.51	16.56/45.27	**24.41/70.43**
70	22.15/57.43	23.03/64.00	22.14/57.39	14.38/32.53	11.23/29.60	14.38/35.96	**23.76/67.07**
80	21.87/55.68	22.56/60.96	21.86/55.66	10.95/15.50	10.07/25.57	11.28/18.87	**23.16/63.84**
90	21.63/54.16	22.12/58.03	21.64/54.16	07.01/03.81	09.04/22.28	07.25/04.69	**22.67/60.89**
100	21.46/52.89	21.69/55.23	21.47/52.92	05.59/01.09	08.13/19.65	05.65/01.27	**22.26/58.21**

In Table V, we compare the PSNR and SSIM of various denoising algorithms on different natural images, with the noise fixed at $\sigma = 50$. From these results, we see that PNLM consistently outperforms NLM in terms of PSNR and SSIM at all noise levels. For a visual comparison, the denoised images obtained using NLM and PNLM are provided in Figure 5. Notice that the output of PNLM is visibly sharper than that obtained using NLM.

For completeness, we have also compared PNLM with the state-of-the-art denoising methods KSVD [16] and BM3D [17]. A couple of representative results are provided in Figures 6 and 7. Notice that the PNLM has successfully reduced the gap in PSNR/SSIM between NLM and K-SVD. This is interesting given that K-SVD is much more complex compared to PNLM. However, there is still a significant margin in the denoising performance with respect to BM3D. In this regard, we wish to make a couple of remarks. First, notice that the denoising using PNLM appears to be somewhat

TABLE IV

COMPARISON OF PSNR/SSIM WITH NLM AND ITS VARIANTS ON THE **HILL** IMAGE [41]. THE HIGHEST PSNRS AND SSIMS ARE MARKED IN BOLD.

σ	NLM	NLEM	PCA-NLM	Zero-CPW	One-CPW	Max-CPW	PNLM
5	34.90/95.30	34.98/96.17	34.66/94.99	28.86/83.41	29.37/84.20	29.06/83.76	**35.78/98.52**
10	30.36/86.55	30.84/89.36	30.23/86.29	28.87/83.57	29.55/84.77	29.07/83.93	**32.15/94.01**
20	26.82/73.85	27.76/81.17	26.79/73.72	28.71/83.76	28.45/87.15	28.90/84.24	**29.21/85.64**
30	25.42/66.20	26.63/76.52	25.40/66.10	27.55/81.89	19.86/64.95	27.42/82.25	**27.67/79.93**
40	24.63/61.40	25.88/72.64	24.61/61.32	24.35/73.41	16.12/47.77	23.51/71.95	**26.52/74.72**
50	24.11/58.21	25.24/68.98	24.10/58.13	20.34/57.32	14.16/38.63	19.48/56.09	**25.63/69.82**
60	23.75/55.99	24.62 /65.58	23.74/55.92	17.05/41.97	12.57/31.87	16.59/43.52	**24.92/65.56**
70	23.48/54.36	24.07/62.37	23.48/54.31	14.62/30.69	11.23/26.50	14.53/26.50	**24.37/61.81**
80	23.27/53.04	23.52/59.36	23.27/52.99	11.65/15.10	10.07/22.33	11.94/18.98	**23.92/58.77**
90	23.08/51.99	23.01/56.55	23.08/51.94	08.00/04.10	09.04/19.06	08.24/05.23	**23.53/56.13**
100	22.94/51.13	22.53/53.90	22.94/51.10	06.60/1.23	08.12/16.39	06.66/01.47	**23.24/54.14**

TABLE V

COMPARISON OF PSNR/SSIM FOR DIFFERENT NATURAL IMAGES [41], [42] AT $\sigma = 50$. THE IMAGES USED ARE BA: BARBARA, CO: COUPLE, HO: HOUSE, BO: BOAT, M: MAN, LE: LENA, AND BR: BRIDGE.

Image	NLM	NLEM	PCA-NLM	Zero-CPW	One-CPW	Max-CPW	PNLM
Ba	22.64/68.91	24.55/79.24	22.63/68.83	19.94/64.66	14.16/44.40	19.26/62.69	**25.41/82.01**
Co	22.59/56.24	23.80/67.56	22.58/56.14	19.92/59.54	14.16/43.17	19.27/58.92	**24.56/71.09**
Ho	24.10/69.59	25.50/63.20	24.09/69.82	20.58/26.27	14.15/12.49	19.67/23.81	**27.35/74.00**
Bo	22.98/62.26	24.01/70.29	22.97/62.19	20.06/58.43	14.16/40.85	19.38/57.22	**25.22/74.01**
Ma	23.65/62.20	24.88/71.73	23.64/62.13	20.20/58.92	14.16/40.62	19.44/57.67	**25.51/73.41**
Le	25.22/75.02	26.46/79.20	25.22/75.00	20.58/57.56	14.14/35.42	19.63/54.28	**27.37/82.11**
Br	19.27/30.87	20.26/41.66	19.26/30.31	17.94/35.56	14.14/31.58	18.10/41.00	**21.12/47.93**

better in homogeneous regions compared to both K-SVD and BM3D. In particular, we see "ripples" and "blocking" artifacts in the homogeneous regions for the K-SVD and BM3D outputs. In fact, the overall visual quality of the PNLM output appears to be better than the K-SVD output for both *Lena* and *Barbara*, though this fact is not quite captured by the quality metrics. The other remark is that the BM3D algorithm uses block-matching and aggregation in its processing pipeline. Therefore, the present idea of pruning can in principle be incorporated into BM3D. We wish to investigate this in future work.

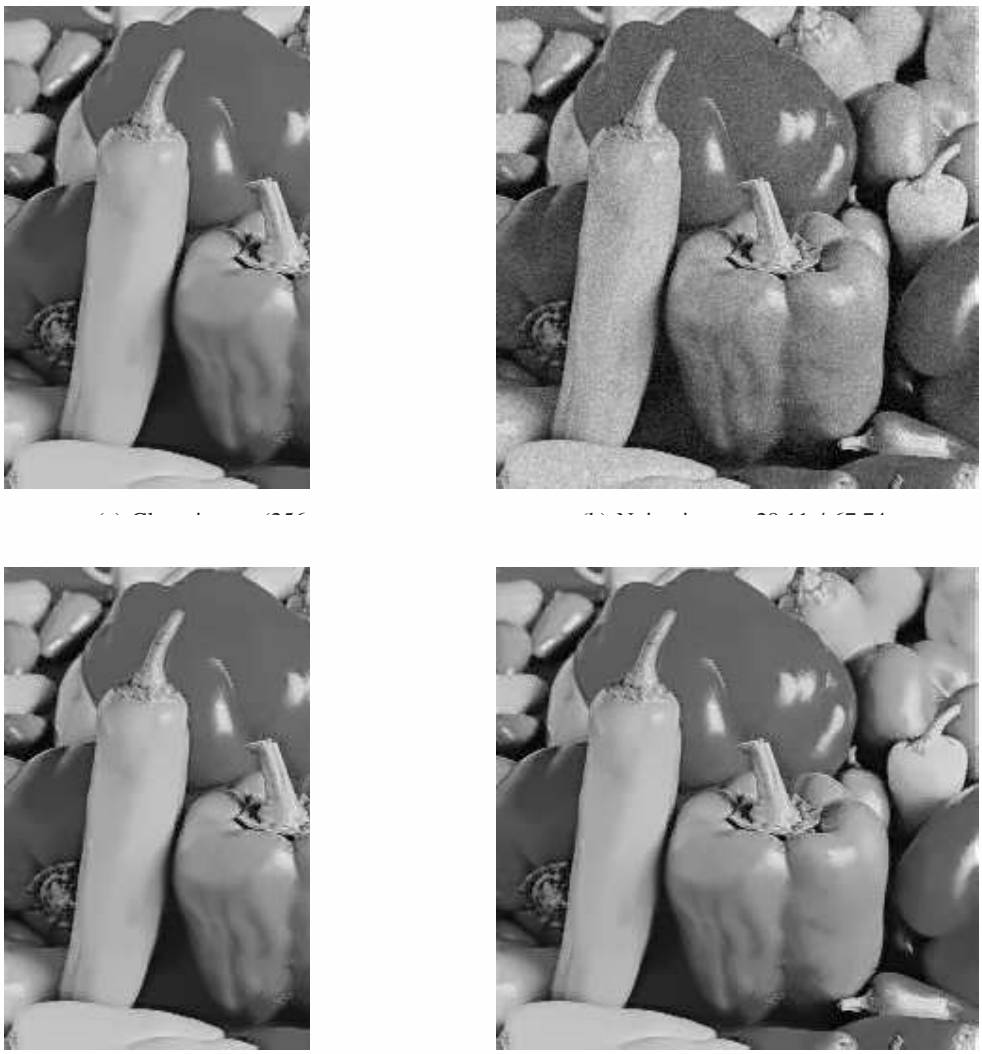

(c) NLM [1]: 32.29 /

(d) PNLM: **33.01 / 90.85**.

Fig. 5. Comparison of the denoising results at $\sigma = 10$ for the *Peppers* image [41]. The corresponding PSNR/SSIM values are shown in the caption.

IV. Conclusion

We demonstrated that the denoising capability of NLM can be significantly improved using pruning, where the pruning is controlled using a single global threshold. The added cost of optimizing the threshold was marginal in comparison to the run-time of NLM. Moreover, we provided several denoising results to demonstrate that pruned NLM outperforms several variants of NLM in terms of PSNR and SSIM. The surprising observation is that, by optimizing a single parameter λ, we could obtain denoising results that are consistently better than that obtained using sophisticated optimization in [33], [34]. It would be interesting to see if the present ideas can be extended to the recently-proposed Vector Nonlocal Euclidean Median [44].

(a) Clean image (512 × 5

(b) Noisy image: 16.08 / 52.62.

(c) NLM: 23.52 / 73.62

(d) PNLM: **26.45 / 84.10.**

(e) K-SVD [16]: 26.88 / 8

(f) BM3D [17]: 27.95 / 93.40.

Fig. 6. Comparison of the denoising results at $\sigma = 40$ for the **Barbara** image [41]. The PSNR/SSIM values are shown in the caption.

(a) Clean image (512 × 5

(b) Noisy image: 18.60 / 52.90.

(c) NLM: 27.30 / 81.63

(d) PNLM: **29.70 / 87.61.**

(e) K-SVD [16]: 30.30 / 8

(f) BM3D [17]: 31.29 / 95.90.

Fig. 7. Comparison of the denoising results at $\sigma = 30$ for the **Lena** image [41]. The PSNR/SSIM values are shown in the caption.

V. Acknowledgements

The authors wish to thank the reviewers for their useful comments and suggestions. This work was supported by a Startup Grant from the Indian Institute of Science and EMR Grant SERB/F/6047/2016-2017 from the Department of Science and Technology, Government of India.

VI. Appendix

In this section, we provide the main steps in the derivation of (10). First, we write (9) as

$$\hat{x}_i = \frac{\sum_{j\in S(i)} \psi(w_{i,j})\ y_j}{\sum_{j\in S(i)} \psi(w_{i,j})},$$

where $\psi(t)$ is given by (11). Then, by the quotient rule for derivatives and after some simplification, we obtain

$$W_i \frac{\partial \hat{x}_i}{\partial y_i} = \psi(1) + \sum_{j\in S(i)} y_j \frac{\partial \psi(w_{i,j})}{\partial y_i} - \hat{x}_i \sum_{j\in S(i)} \frac{\partial \psi(w_{i,j})}{\partial y_i}, \tag{14}$$

where W_i is given by (11). Applying chain rule, we have

$$\frac{\partial \psi(w_{i,j})}{\partial y_i} = \psi'(w_{i,j}) \frac{\partial w_{i,j}}{\partial y_i}, \tag{15}$$

where

$$\frac{\partial w_{i,j}}{\partial y_i} = \frac{2w_{i,j}}{h^2} \cdot \begin{cases} y_j + y_{2i-j} - 2y_i & i-j \in \mathcal{P} \\ y_j - y_i & i-j \notin \mathcal{P}, \end{cases} \tag{16}$$

and

$$\psi'(t) = \frac{\left(1 + (1+\alpha t)e^{-\alpha(t-\lambda)}\right)}{(1+e^{-\alpha(t-\lambda)})^2} \tag{17}$$

Substituting (15) and (16) in (14),

$$\begin{aligned} W_i \frac{\partial \hat{x}_i}{\partial y_i} =& \Big\{ \psi(w_{i,i}) + \sum_{j\in S(i)} \psi'(w_{i,j})(y_j - \hat{x}_i) \frac{\partial w_{i,j}}{\partial y_i} \Big\} \\ =& \Big\{ \psi(w_{i,i}) + \sum_{j\in S(i)} \frac{2w_{i,j}\ \psi'(w_{i,j})}{h^2} (y_j - \hat{x}_i)(y_j - y_i) \\ &+ \sum_{k\in\mathcal{P}} \frac{2w_{i,i+k}\ \psi'(w_{i,i+k})}{h^2} (y_{i+k} - \hat{x}_i)(y_{i-k} - y_i) \Big\}. \end{aligned}$$

After performing some reductions, we arrive at (10):

$$\begin{aligned} \frac{\partial \hat{x}_i}{\partial y_i} =& \frac{1}{W_i} \Big\{ \psi(1) + \sum_{j\in S(i)} \frac{2w_{i,j}\ \psi'(w_{i,j})}{h^2} (y_j - \hat{x}_i)(y_j - y_i) \\ &+ \sum_{k\in\mathcal{P}} \frac{2w_{i,i+k}\ \psi'(w_{i,i+k})}{h^2} (y_{i+k} - \hat{x}_i)(y_{i-k} - y_i) \Big\}. \end{aligned}$$

REFERENCES

[1] Buades, A., Coll, B., and Morel, J. M.: 'A non-local algorithm for image denoising,' *Proc. IEEE International Conference on Computer Vision and Pattern Recognition,* Jun 2005, vol. 2, pp. 60-65.

[2] Buades, A., Coll, B., and Morel, J. M.: 'A review of image denoising algorithms, with a new one,' *Multiscale Modeling and Simulation*, vol. 4, pp. 490-530, 2005.

[3] Salmon, J.: 'On two parameters for denoising with non-local means,' *IEEE Signal Processing Letters,* 2010, **17**, (3), pp. 269-272.

[4] Duval, V., Aujol, J., and Gousseau, Y.: 'A bias-variance approach for the nonlocal means,' *SIAM Journal on Imaging Sciences*, 2011, **4**, (2), pp. 760-788.

[5] Wu, Y., Tracey, B., Natarajan, P., and Noonan, J. P.: 'James-Stein type center pixel weights for non-local means image denoising,' *IEEE Signal Processing Letters,* 2013, **20**, (4), pp. 411-414.

[6] Feng, J., Li Song, *et al.*: 'An optimized pixel-wise weighting approach for patch-based image denoising,' *IEEE Signal Processing Letters*, 2015, **22**, (1), 115-119.

[7] Zhang, X., Feng, X., and Wang, W.: 'Two-direction nonlocal model for image denoising,' *IEEE Transactions on Image Processing,* 2013, **22**, (1), pp. 408-412.

[8] Liu, G., Zhong, H., and Jiao, L.: 'Comparing noisy patches for image denoising: a double noise similarity model,' *IEEE Transactions on Image Processing,* 2015, **24**, (3), pp. 862-872.

[9] Rehman A., and Wang, Z.: 'SSIM-based non-local means image denoising,' *Proc. IEEE International Conference on Image Processing*, 2011, pp. 217-220.

[10] Orchard, J., Ebrahimi, M., and Wong, A.: 'Efficient non-local means denoising using the SVD,' *Proc. IEEE International Conference on Image Processing*, 2008, pp. 1732-1735.

[11] Tasdizen, T.: 'Principal neighborhood dictionaries for non-local means image denoising,' *IEEE Transactions on Image Processing*, 2009, **18**, 12, pp. 2649-2660.

[12] Van De Ville, D., and Kocher, M.: 'Non-local means with dimensionality reduction and SURE-based parameter selection,' *IEEE Transactions on Image Processing*, 2011, **20**, (9), pp. 2683-2690.

[13] Zuo, Chenglin, et al.: 'Image denoising using quadtree-based nonlocal means with locally adaptive principal component analysis,' *IEEE Signal Processing Letters*, 2016, **23**, (4), pp. 434-438.

[14] Tracey, B. H., and Miller, E. L.: 'Nonlocal means denoising of ECG signals,' *IEEE Transactions on Biomedical Engineering,* 2012, **59**, (9), pp. 2383-2386.

[15] Yadav, S. K., Sinha, R., and Bora, P. K.: 'Electrocardiogram signal denoising using non-local wavelet transform domain filtering,' *IET Signal Processing,* 2015, **9**, (1), pp. 88-96.

[16] Elad, M., and Aharon, M.: 'Image denoising via sparse and redundant representations over learned dictionaries,' *IEEE Transactions on Image Processing,* 2006, **15**, (12), pp. 3736-3745.

[17] Dabov, K., et al.: 'Image denoising by sparse 3D transform-domain collaborative filtering,' *IEEE Transactions on Image Processing,* 2007, **16**, (8), pp. 2080-2095.

[18] Hossein, T., and Milanfar, P.: 'Global image denoising,' *IEEE Transactions on Image Processing*, 2014, **23**, (2), pp. 755-768.

[19] Duan, J., et al.: 'Second order Mumford-Shah model for image denoising,' *Proc. IEEE International Conference on Image Processing*, 2015.

[20] Yue, H., et al.: 'Image denoising by exploring external and internal correlations,' *IEEE Transactions on Image Processing*, 2015, **24**, (6), pp. 1967-1982.

[21] Romano, Y., and Elad, M.: 'Boosting of image denoising algorithms,' *SIAM Journal on Imaging Sciences,* 2015, **8**, (2), pp. 1187-1219.

[22] Dinh, K. Q., Canh, T. N., and Byeungwoo, J.: 'Color Image Denoising via Cross-Channel Texture Transferring,' *IEEE Transactions on Image Processing,*, 2016, **23**, (8), pp. 1071-1075.

[23] Chen, Yang, et al.: 'Improving low-dose abdominal CT images by weighted intensity averaging over large-scale neighborhoods,' *European Journal of Radiology,* 2011, **80**, (2), e42-e49.

[24] Chen, Yang, et al.: 'Thoracic low-dose CT image processing using an artifact suppressed large-scale nonlocal means,' *Physics in Medicine and Biology,* 2012, **57**, (9), pp. 2667-2688.

[25] Chen, F., Xiaojun H., and Chen, W.: 'Texture-preserving image deblurring,' *IEEE Signal Processing Letters,* 2010, **17**, (12), pp. 1018-1021.

[26] Jung, M., et al.: 'Nonlocal Mumford-Shah regularizers for color image restoration,' *IEEE Transactions on Image Processing,* 2011, **20**, (6), pp. 1583-1598.

[27] Wong, A., and Orchard, J.: 'A nonlocal-means approach to exemplar-based inpainting,' *Proc. IEEE International Conference on Image Processing,* 2008.

[28] Buades, A., et al.: 'Self-similarity driven color demosaicking,' *IEEE Transactions on Image Processing,* 2009, **18**, (6), pp. 1192-1202.

[29] Protter, M., et al.: 'Generalizing the nonlocal-means to super-resolution reconstruction,' *IEEE Transactions on Image Processing,* 2009, **18**, (1), pp. 36-51.

[30] Zhang, K., et al.: 'Single image super-resolution with non-local means and steering kernel regression,' *IEEE Transactions on Image Processing,* 2012, **21**, (11), pp. 4544-4556.

[31] Zhang, K., et al.: 'Learning multiple linear mappings for efficient single image super-resolution,' *IEEE Transactions on Image Processing,* 2015, **24**, (3), pp. 846-861.

[32] Deledalle, C.-A., Duval, V., and Salmon, J.: 'Non-local methods with shape-adaptive patches (NLM-SAP),' *Journal of Mathematical Imaging Vision,* 2012, **43**, (2), pp. 103-120.

[33] Chaudhury, K. N., and Singer, A.: 'Non-local Euclidean medians,' *IEEE Signal Processing Letters*, 2012, **19**, (11), pp. 745-748.

[34] Sun, Z., and Chen, S.: 'Analysis of non-local Euclidean medians and its improvement,' *IEEE Signal Processing Letters*, 2013, **20**, (4), pp. 303-306.

[35] Chaudhury, K. N., and Singer, A.: 'Non-local patch regression: Robust image denoising in patch space,' *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2013, pp. 1345 - 1349.

[36] Stein, C.: 'Estimation of the mean of a multivariate normal distribution,' *The Annals of Statistics,* 1981, **9**, pp. 1135-1151.

[37] Van De Ville, D., and Kocher, M.: 'Sure-based non-local means,' *IEEE Signal Processing Letters*, 2009, **16**, (11), pp. 973-976.

[38] Kiefer, J.: 'Sequential minimax search for a maximum,' *Proc. American Mathematical Society*, 1953, **4**, (3), pp. 502-506.

[39] Ghosh, S., and Chaudhury, K. N.: 'Fast Separable Non-Local Means,' *SPIE Journal of Electronic Imaging,* 2016, **25**, (2), 023026.

[40] Darbon, J., Cunha, A., *et al.*: 'Fast nonlocal filtering applied to electron cryomicroscopy,' *Proc. IEEE International Symposium on Biomedical Imaging*, 2008, pp. 1331-1334.

[41] 'BM3D Image Database', http://www.cs.tut.fi/~foi/GCF-BM3D.

[42] 'Computer Vision Test Images', http://www.cs.cmu.edu/~cil/v-images.html.

[43] Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P.: 'Image quality assessment: From error visibility to structural similarity,' *IEEE Transactions on Image Processing*, 2004, **13**, (4), pp. 600-612.

[44] Lin C.-Y., Minasian A., Qi X.-J., and Wu H.-T.: 'Vector nonlocal Euclidean median: principal bundle captures the nature of patch space,' *arXiv:1611.05073*, 2016.